

CLAIMS

What is claimed is:

- 1 1. A system for exchanging routing information in one or more networks,
2 the one or more networks including a plurality of at least partially
3 interconnected nodes, the protocol comprising:
4
5 a plurality of path vectors for routes in the one or more networks, the
6 plurality of path vectors included in the routing information;
7
8 a multi-tier hierarchy amongst the plurality of nodes in the one or more
9 networks, such that the one or more networks are operative to expand or
10 summarize the routing information to select nodes in the plurality of nodes
11 based on a rank of the select nodes in the multi-tier hierarchy;
12
13 a flooding mechanism for exchanging the routing information amongst
14 the plurality of nodes;
15
16 a link-state database in each of the plurality of nodes, the link state
17 database including a virtual topology of the one or more networks, such that
18 each of the plurality of nodes is operative to generate the link state database
19 from the routing information, the link-state database further including the
20 plurality of path vectors for routes in the one or more networks.

- 1 2. The system of claim 1, wherein a convergence time of the one or more
2 networks exchanging the routing information via the protocol is less than
3 an average convergence time for a topologically equivalent network
4 connected via OSPF.

- 1 3. The system of claim 1, wherein a convergence time of the one or more networks exchanging the routing information via the protocol is less than an average convergence time for a topologically equivalent network connected via BGP.
- 1 4. The system of claim 1, wherein the one or more networks includes one or more autonomous systems.
- 1 5. The system of claim 4, wherein the one or more networks includes two or more autonomous systems.
- 1 6. The system of claim 5, wherein each of the plurality of nodes maintains a list of logically adjacent nodes from the plurality of nodes.
- 1 7. The system of claim 6, wherein the list of logically adjacent nodes are non-equivalent to physically adjacent nodes.
- 1 8. The system of claim 7, wherein two or more logically adjacent nodes from the plurality of nodes reside on two or more distinct autonomous systems from the one or more networks.
- 1 9. The system of claim 1, wherein each of the plurality of nodes is operative to populate the link-state database from a shortest path first algorithm.
- 1 10. The system of claim 9, wherein the shortest path first algorithm is a modified Dijkstra algorithm.
- 1 11. The system of claim 1, wherein each of the plurality of nodes is operative to create adjacencies other nodes in the one or more networks via a four-way handshake.

- 1 12. The system of claim 11, wherein the protocol includes a hello message,
- 2 such that the hello message is exchanged periodically between
- 3 adjacent nodes after the four-way handshake.
- 1 13. The system of claim 12, wherein the hello message includes a modified
- 2 hello PDU with one or more additional parameters.
- 1 14. The system of claim 1, wherein the multi-tier hierarchy includes one or
- 2 more higher level tiers, such that nodes in the one or more higher level
- 3 tiers are in communication via an Exterior Gateway Protocol (EGP).
- 1 15. The protocol of claim 14, wherein the EGP is a version of Border
- 2 Gateway Protocol.
- 1 16. The protocol of claim 1, wherein the multi-tier hierarchy includes one or
- 2 more lower level tiers, such that nodes in the one or more lower level
- 3 tiers are in communication via an Interior Gateway Protocol (IGP).
- 1 17. The protocol of claim 16, wherein the IGP is a link state protocol.
- 1 18. The protocol of claim 17, wherein the IGP is one of OSPF and IS-IS.
- 1 19. A method of selecting routes at a first node in a communications
- 2 network, the method comprising:
 - 3 establishing a plurality of nodes logically adjacent to the first node,
 - 4 establishing the plurality of nodes further including completing a four way
 - 5 handshake with each of the plurality of logically adjacent nodes;
 - 6 receiving a plurality of routing tables at periodic intervals from the
 - 7 plurality of adjacent nodes;

8 populating a routing table local to the first node, populating the local
9 routing table further including selecting a plurality of routes to the plurality of
10 nodes from the routing tables, selecting the plurality of routes further including
11 determining a path length for each of the plurality of routes and applying a
12 policy vector to each of the plurality of routes, applying the policy vector
13 including generating one or more metrics for discriminating between the
14 plurality of routes.

1 20. The method of claim 19, wherein the one or more metrics are in a
2 prioritized order.

1 21. The method of claim 19, wherein the selecting the plurality of routes
2 further includes resolving ties between two or more routes in the
3 plurality of routes.

1 22. The method of claim 21, wherein the path length for the two or more
2 routes are identical.

1 23. The method of claim 22, wherein resolving ties between the two or more
2 routes further includes selecting a route from the two or more routes
3 based on the one or more metrics.

1 24. The method of claim 23, wherein the one or more metrics includes BGP
2 path attributes.

1 25. The method of claim 23, wherein the one or more metrics includes BGP
2 Multi Exit Discriminator attributes.

1 26. The method of claim 23, wherein the one or more metrics includes
2 autonomous system path lengths from the two or more routes.

1 27. The method of claim 19, further comprising:

2 selecting one or more optimal routes from the plurality of routes based
3 on the one or more metrics.

1 28. The method of claim 27, wherein the one or more optimal routes have
2 minimal values for the one or more metrics.

1 29. The method of claim 27, wherein the one or more optimal routes ensure
2 that the communications network is load balanced.

1 30. The method of claim 27, wherein the one or more optimal routes have a
2 minimal length.

1 31. The method of claim 27, wherein the one or more metrics includes a
2 distance metric indicating, for each of the two or more routes, a length
3 of an internal gateway path traversed by the two or more routes.